LLNL-TR-416510

LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

NNSA'S Computing Strategy, Acquisition Plan, and Basis For Computing Time Allocation

Daniel J. Nikkel

August 10, 2009

Report Mandated by the 2009 ASC Appropriation
(H.R. 1105; Public Law 111-8)

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Contents

1 Executive Summary

This report is in response to the Omnibus Appropriations Act, 2009 (H.R. 1105; Public Law 111-8) in its funding of the National Nuclear Security Administration's (NNSA) Advanced Simulation and Computing (ASC) Program. This bill called for a report on ASC's plans for computing and platform acquisition strategy in support of stockpile stewardship. The appropriations bill called for this report to address three specific issues:

1. *Identify how computing capability at each of the labs will specifically contribute to stockpile stewardship goals, and on what basis computing time will be allocated to achieve the goal of a balanced program among the labs.*

Capability computing is supporting all major stockpile stewardship mission areas, and through a proposal process all three laboratories have received allocations of time necessary to address mission priorities, as seen in the data summarized in the Appendix in Section 10. The specific activities and computing needs at each laboratory to support ongoing stockpile stewardship, which vary from quarter to quarter, can be loosely divided into three major areas: Directed Stockpile Work (DSW), support of experiments and the advance towards predictive capability. DSW includes Life Extension Programs (LEPs), Significant Finding Investigations (SFIs), annual assessments and evaluation of stockpile options. The nature of the work both across and within these three areas can present different computing requirements. In fact, each activity tends to concentrate on particular computing regimes; *predominant* regimes for each area are indicated in the table below.

		Computing Regimes			
		High throughput	High-resolution multi-physics	High-resolution weapons science	UQ studies
SSP mission needs	DSW	X	X		X
	Predictive capability		X	X	X
	Support of Experiments	X	X		

Table 1.1– Predominant regimes of computing used in support of SSP mission needs.

ASC's platform strategy to cost effectively address the various regimes of computing involves three types of major systems—*capacity*, *capability*, and *advanced architecture*s.

- Low-cost *capacity* systems, which are provided to all three laboratories, generally represent lower-risk, less expensive production computing systems that run parallel problems with more modest (but nonetheless significant on an absolute scale) computational requirements. These are used to address the need for *high-throughput* (see Table 1.1). Computing time on capacity systems is allocated to users locally by each laboratory, and a common tri-lab software environment on these systems enables work to be shifted among laboratories in exigent circumstances and saves money on operations.
- *Capability* systems are remotely available to all three laboratories and are run as NNSA user computing facilities. These systems have the computational power, memory size, and interconnect speed necessary to address the most challenging weapons system computing regimes – in particular, the *high-resolution, multi-physics* regime shown in Table 1.1. Time on capability systems is allocated through a tri-laboratory review process of proposals, with the two primary criteria for access being programmatic priority and need to run in the high resolution, multi-physics regime.
- *Advanced architecture* systems enable the development and exploration of promising, low-power-consuming approaches for future general-purpose capability and capacity systems. These machines also have successfully supported *high-resolution weapons science* simulations necessary to improve the accuracy of critical physics models in the design codes and have also delivered important results (*e.g.*, Pu aging) for the Stockpile Stewardship Program (SSP). Advanced architecture systems have the additional advantage that they engage the major high performance computing (HPC) companies in ways that help steer commercial technology to solutions that can be leveraged for national security needs. Allocation of time on advanced architecture systems in the future will vary from machine to machine, depending on its state of technological maturity; however, it is expected that the use of these machines by all three laboratories will be accommodated within the user facility context.

ASC has adopted a strategy wherein capability and advanced architecture systems will be sited at two national NNSA user facilities (one in California and one in New Mexico) and will be remotely available to all three laboratories. Two user facilities are a lowest cost solution for the program to provide the demanding infrastructure required by these systems while mitigating the risk of a single-point-of-failure, which could completely deprive the complex of the most essential computing resources for many months or years and impact national security. The ASC Program has developed a strategy of employing the appropriate computer for the appropriate calculation in order to most efficiently utilize expensive resources. While capacity systems are the utility computers for stockpile work across the tri-laboratory, capability and advanced architecture systems are acquired to support the most demanding SSP goals. As such, the mission need for capability and advanced architecture systems are documented in formal project management critical decisions. [See Sections 3 and 6 for a more expanded discussion, and Section 10 for a summary of data on capability computing allocation and application to SSP problems.]

2. *Explain the NNSA's acquisition strategy for capacity and capability of machines at each of the labs and how it will fit within the existing budget constraints*.

ASC's platform acquisition principles and broad acquisition plan for computing hardware are described in the *ASC Platform Strategy*[*] document. ASC acquires machines to provide computing resources to the SSP as a whole, based on mission needs, not by considering each lab separately. Each year, the ASC Program evaluates its ability to accomplish its mission and prioritizes its work scope within available resources across a five-year horizon. A detailed, budget-driven plan for ASC's acquisition of the three types of major systems is laid out over the next five years based on prevailing guidance of future budgets and is updated regularly.

The current five-year plan provides for the siting of capability and advanced architecture systems at both NNSA user facilities to most effectively use both facilities to support the needs of the program. The plan currently provides for a new capability and a new advanced architecture system approximately every four years. As was described in the response to issue #1, all three laboratories will use these systems with the criteria for access being programmatic need and computing regime required. The plan also calls for replacing lower-cost capacity resources at all three laboratories on a three-year cycle. A single procurement is done for all three laboratories in order to accrue economies of scale. In the event that budgets decrease, ASC HQ can make a decision to either (1) extend the interval between procurements or (2) cancel procurements. The particulars of the decision are based on the current exigencies, including among other factors, the character of contracts in place and the criticality of systems to mission need and the impacts of their loss at the time of the budget adjustment.

The acquisition plan is currently based on the assumption of a flat funding profile and will maintain, and to some degree extend, the current computational capability of the program. However, delivering exascale computing on a time-scale synchronized with the *ASC Roadmap*[**] will require an initiative beyond the current ASC program and indeed beyond those of any currently existing government program. To tackle this problem ASC and the DOE Office of Science Advanced Scientific Computing Research (ASCR) program have announced an exascale collaboration, the goal of which is commercially viable exascale computing for scientific applications by the end of the next decade. Such a goal is beyond the wherewithal of even the combined programs. Consequently, as the collaboration goals translate into an executable plan, stakeholders will be briefed on resource requirements and deliverables. [See Sections 4.2, 5, and 7 for a more expanded discussion.]

---

[*] *A Platform Strategy for the Advanced Simulation and Computing Program*, Robert Meisner, August 2007, NA-ASC-113R-07-Vol. 1-Rev. 0.

[**] *Advanced Simulation & Computing Roadmap: National Nuclear Security Through Leadership in Weapons Science*, Dimitri F. Kusnezov and Njema Frazier, November 2006, NA-ASC-105R-06-Vol. 1-Rev. 0.

3. *Identify the technical challenges facing the program and a strategy to resolve them*.

As the stockpile ages and the state of weapons drifts further from the events in the historical test base, the need for a predictive computational capability increases inexorably. The SSP's needs for the next decade therefore compel ASC to develop increasingly predictive codes that require dramatically more powerful computer platforms than are currently available. In short, two broad challenges face the ASC Program: (1) establishing the science base to enable certification of the future stockpile in the absence of nuclear testing—that is, (1a) developing improved physics and engineering models that will be incorporated into the design codes so that, (1b) within an uncertainty quantification (UQ) discipline, the program can *predict* with confidence a weapon's performance, safety, and reliability (this is loosely called *predictive simulation* or *predictive capability*); and (2) delivering the technologically advanced (exascale) environment required by predictive simulation. The latter includes (2a) hardware platforms, (2b) software environments, (2c) algorithm development, and (2d) integrated design codes compatible with the exascale programming environments.

The first of these challenges requires collaboration between elements of NNSA, including ASC, the science and engineering campaigns, and Directed Stockpile Work (DSW). The NNSA laboratories are actively engaged in science and computing programs to reduce the reliance on calibration over the next decade in order to develop a more predictive capability to support the stockpile and to support decisions and options anticipated for new requirements such as safety and surety.

For the second of these challenges, the natural evolution of existing computer technologies will not carry the program to exascale. Achieving exascale computing will require a funded DOE collaboration over the next decade that engages other government agencies, industry and academia. As a consequence, NNSA laboratories (with ASC HQ) are working closely with Office of Science laboratories (with ASCR HQ) to develop a comprehensive plan for a DOE initiative for exascale simulation for open science and national security. As plans coalesce they will be presented to stakeholders including the key congressional committees.

[See sections 3 and 4 for an expanded discussion of these challenges, and ASC's strategy to address them.]

This report was submitted for review to an independent panel of experts from academia and the DOE Office of Science laboratories familiar with HPC and the ASC Program. The reviewers are listed in Section 9.

2 Congressional Report Request

This report was produced in response to a requirement in the 2009 funding of the ASC program in the Omnibus Appropriations Act, 2009 (H.R. 1105; Public Law 111-8). In particular, this bill states:

> The bill provides $556,125,000 for Advanced Simulation and Computing, including $15,000,000 to develop the new Zia Platform. The budget submitted by NNSA has a striking lack of detail regarding the NNSA's computing strategy, acquisition plan, and on what basis computing time will be allocated among the national labs. This raises the concern that the acquisition strategy for new platforms will not fit within the available budget. As computing is an essential tool in stewardship, the NNSA is directed to provide a written report addressing the following issues: 1) identify how computing capability at each of the labs will specifically contribute to stockpile stewardship goals, and on what basis computing time will be allocated to achieve the goal of a balanced program among the labs; 2) explain the NNSA's acquisition strategy for capacity and capability of machines at each of the labs and how it will fit within the existing budget constraints; and 3) identify the technical challenges facing the program and a strategy to resolve them. This report shall have the benefit of independent review, and be submitted to the House and Senate Appropriations Committees within 6 months after enactment of this Act.

This report begins by describing the role of ASC in the SSP. It then describes the major technical challenges facing ASC in supporting the SSP in the coming years along with a strategy to address them. Following that, it explains the platform acquisition strategy to supply the needed computing resources to the SSP, and the way that computing time on the various types of machines is allocated.

The three specific issues the bill requires to be addressed by this report are responded to briefly in the Executive Summary of Section 1, but are expanded upon in the subsequent document. Issue #1, regarding how computing capability contributes to SSP goals and how computing time is allocated, is addressed in Sections 3 and 6, as well as by the data in Section 10. Issue #2, regarding the acquisition strategy for computing platforms, is addressed in Sections 5, 4.2, and 7. Issue #3, regarding the challenges facing the program and the strategy to address them, is addressed in Sections 3 and 4.

3 Introduction

3.1 The Science-Based Stockpile Stewardship Program and ASC

The science-based SSP is an integrated program involving design analysis, stockpile surveillance, integrated and focused physics experiments, development of improved theoretical understanding of weapons physics and engineering, and simulation science. In the absence of nuclear testing, the foundation for maintaining the stockpile is computational simulation that is informed and validated by archival data and the science and engineering campaigns. The ASC Program develops for and provides to the SSP the computational hardware, software environments, application codes, theoretical models, and validation processes to underpin the use of simulations with confidence in assessing the current stockpile as well as future stockpile options. ASC's programmatic success results from it being a balanced program, driven by SSP mission deliverables, that provides all the elements necessary for a simulation capability, not just hardware. This mission drive and balance has also differentiated it from many other HPC efforts. ASC has, since its inception, been driven by the need to ensure the safety, reliability, and performance of the nuclear weapons stockpile without nuclear testing.[1] The ASC Program mission, strategies, goals, and execution process are presented in a group of documents produced by the NNSA.[1–4]

The program, which began at about the same time as the moratorium on nuclear testing, has had two major phases: the Accelerated Strategic Computing Initiative (ASCI) and the ASC Program. ASCI focused on developing three-dimensional (3D) tools and the machines capable of running those tools. The goal, based on what was viewed as the initial entry-level computational capability to enable stockpile stewardship without nuclear testing, was to develop computer platforms capable of 100 teraFLOPS (~0.1 petaFLOPS) performance, along with science and 3D integrated design codes capable of scaling to this level. This goal represented a 1000-fold increase in computing power from program inception. It was achieved with a new generation of design codes, software tools, and the computational resource of the Purple machine in 2005 and represented an extremely difficult technical achievement in both hardware and software, requiring effective teaming between NNSA, the labs, and industry. The solution involved developing simulation codes using a programming model based on the standardized Message Passing Interface (MPI) to be run on platforms assembled by connecting many processors together through high-speed interconnects and taking advantage of weak scaling (problem size scaling linearly in the number of processors). Continuing this MPI-based programming model, computing power has been extended to the petascale regime, but this is near the limit for MPI-only simulations. Integrated design codes have been developed to scale to petascale machines using the MPI programming model.

There are key elements of physics in these codes that are not sufficiently well understood, and models with incomplete physics are utilized requiring stockpile simulations to be *calibrated* with historical nuclear test data using adjustable parameters (loosely referred to as "knobs"). This paradigm faces a major limitation going into the future as the state of weapons diverge from their "as-tested" condition in the historic test base while each weapon within a particular class also ages away from other systems in its class.

As the new computing capabilities and tools became essential elements of ongoing work in the SSP, the ASCI initiative to develop the entry-level computing capability and tools transitioned to the ASC Program.[1] ASC is characterized by the mandate to develop a predictive capability in our integrated design codes. With predictive codes (used with appropriate UQ methodologies), simulations of systems diverging from the historical test database could be made with confidence, as the integrated design codes would no longer be calibrated to the test database. The historical nuclear test data will still have great value, as it will be used for *validation* of capabilities, not *calibration* of simulations.

3.2 The Role of Computing in Stockpile Stewardship

As computational simulation has become an essential component of stockpile stewardship, the capabilities enabled by ASC-class computing have put computational science on an equal footing with theoretical and experimental science as a tool for studying basic issues of weapons science and for scientific discovery. ASC platforms and tools are also the primary computational resources for ongoing, time-constrained work on the stockpile. The two major foci of the ASC Program strategy are: (1) to meet the continuing and time-constrained needs of stockpile stewardship; and (2) to ensure progress toward the long-term goal of reduced dependence on phenomenology to enhance confidence (i.e., to move from calibrated simulations to predictive simulations).[1]

Under the first focus area, meeting the continuing needs of stockpile stewardship, ASC simulations using ASC codes that are run on ASC computers are heavily used by DSW in supporting LEPs and SFIs and for the annual assessment of systems in the stockpile. Responding to DSW imperatives is a major requirements driver for the ASC platform acquisition timetable. ASC simulations also support SSP experimental programs, experimental design (e.g., Dual-Axis Radiographic Hydrodynamic Test [DARHT] experiments, targets for the National Ignition Facility [NIF]), and the analysis of results for integrated experiments like large hydrodynamic tests. In addition to the integrated design codes that are used directly on stockpile system simulations, ASC develops specialized physics and material property codes that are used to study basic issues of weapons science and develop models and data used by the integrated design codes.

As the complex systems in the stockpile continue to age, we need the ability to simulate the evolution of the state of the weapons and predict their performance in their current state. The second focus of the ASC Program therefore addresses the need to develop a predictive simulation capability with improved physical models as stockpile systems age further away from the historical test data. In the absence of nuclear testing, confidence in the simulations through quantifiable measures of the uncertainty in the predictions becomes essential. As was stated in the *ASC Roadmap*[3], "… sustaining the testing moratorium requires that we transition to a point of sustainability at which our confidence in science-based simulations exceeds our confidence in simulations calibrated by underground test data." Whatever the future of the stockpile, it will be essentially different from what it is today (older, smaller, less diverse in systems but more diverse in individual differentiation), and future decisions regarding the stockpile will rely critically on ASC simulation capabilities in assuring the certainty to friend and foe in the readiness of the US deterrent.

3.3 Stockpile Decisions and Maintenance Options

In addition to overseeing stewardship and assessment of the current stockpile, the NNSA counsels policymakers on options for the stockpile of the future. The current stockpile is aging, and periodic changes need to be made to keep it safe, secure, and reliable. To date these changes have been effected through LEPs, which are basically a refurbishment of an existing system. Whether or not the stockpile is smaller or less diverse in future, the SSP needs to preserve the capability to evaluate a spectrum of options, often involving safety or surety, regarding the future stockpile. Such decisions are more challenging in an environment where full-system nuclear testing is not available. Given the nation's self-imposed moratorium on nuclear testing and the possibility of moving to an era with a Comprehensive Test Ban Treaty (CTBT), a robust, predictive simulation capability will become even more critical to stockpile stewardship as the value of each weapon in a limited stockpile is obviously enhanced. Evaluation of these stockpile options will require a new kind of computational approach—one that does not depend on calibration to an increasingly irrelevant test base.

3.4 Supporting Broader National Security Missions

NNSA, through ASC, has for more than a decade been at the forefront of HPC in development of computer platforms and computational and scientific tools to support the stockpile stewardship mission. These capabilities are not, however, all limited in application to stockpile-related problems. Applications of ASC computational tools and capabilities benefit broader national security missions and other problems of national interest, and as the ASC Program matures, there will be increasing dependence on its capabilities from the larger national security community. This represents both an opportunity and a challenge for the program. Areas outside stockpile stewardship for which ASC capabilities and tools are applied include nuclear forensics, nuclear counterterrorism, seismic modeling for nonproliferation, radiation hardening and survivability for microelectronics, vulnerabilities of critical infrastructure, weapon effects, and foreign assessments. ASC resources have also been used in the space shuttle Columbia investigation and in missile defense simulations (e.g., Operation Burnt Frost), which demonstrate the agility of the program in responding to exigent national needs.

4 Technical Grand Challenges Facing the ASC Program

4.1 Enabling SSP Certification of the Future Stockpile

Simulation codes currently used for stockpile applications involve multiple complex physics models. Within these codes are a number of ad hoc physically incomplete models with "knobs" or tunable parameters. These "knobs" are set so that the simulations match a number of past full-system nuclear tests (but not all). However, different classes of systems often require different settings, and the degree of confidence in the code's ability to predict diminishes as the weapons age. As time passes or changes are made to meet DSW requirements, the state of the systems drift ever further from the conditions they were in when they were tested. Consequently, simulations that rely on *calibration* against test data become less reliable. *Over time, just maintaining the status quo in code capabilities will allow uncertainties in assessing the stockpile to increase*. Because of this, it is imperative to develop a *predictive* simulation capability—that is, the ability to predict with quantified uncertainty (and without recourse to calibration) how a weapon will perform. To this end, the NNSA SSP has developed an integrated roadmap known as the Predictive Capability Framework (PCF). Cutting across program elements, the PCF sets significant peg posts for progress over the next decade. For ASC, these peg posts relate to the replacement of several major ad hoc "knobs" in the current simulation capability with increasingly predictive physics-based models. The major improved models include energy balance, boost and secondary performance, as seen in the figure at the end of Section 4.1.2. Moreover, as assessment of the stockpile becomes more simulation based, an essential component of this roadmap is to calculate, measure, and understand the uncertainty in the predictions. Achieving this predictive simulation capability requires a balanced approach of scientific model development using theory and experiment, improved computer science algorithms, hardware capable of running the more complex multi-physics codes at sufficient resolution to capture essential phenomena and the rigorous application of uncertainty quantification technology.

4.1.1 Providing Stockpile Decision and Maintenance Options

In making decisions whether to modernize the stockpile or move to a smaller or less diverse stockpile, policymakers need to understand the impact of various decisions regarding safety, security, and reliability. Issues will include, for example, investigating the use of new materials to replace materials that are no longer manufactured or that could make systems safer. Answering such questions without full-scale testing will require the more predictive models under development.

As mentioned earlier, meeting this goal requires coordinated efforts from the model development and implementation elements of ASC as well as the experimental elements of the NNSA science and engineering campaigns. Sustained support in the science and model development areas will be needed to make sufficient progress to meet stockpile program and DSW requirements.

In addition to improved physics fidelity, another factor limited by computational power is geometric fidelity (i.e., resolution). Standard design calculations run with integrated codes are typically not numerically converged. Running the same problems with finer resolution on more powerful computers will improve results and reduce uncertainties. Highly resolved simulations enabled by ASC-class machines have revealed phenomena not captured in under-resolved simulations. It is in this realm of scientific discovery with adequately resolved simulations that the key weapons science issues facing the program will have to be explored.[5]

Finally, reliance on simulation requires assurance that simulations are sufficiently accurate for making decisions: this is achieved through quantitative assessment of uncertainties. A significant component of the ASC Program involves developing the methodologies of validation and UQ that are necessary to ensure such a predictive capability. The large number of large-scale simulations necessary to perform sensitivity studies, an aspect of UQ, requires massive computing throughput.

In summary, predictive simulation depends on advances in the fidelity of physics, the accuracy of numerical methods, and our ability to assess uncertainty—all *ASC Roadmap* goals. These advances are, in turn, dependent on the level of computing that can be brought to bear.

4.1.2 Realizing Predictive Simulation Capability through Exascale Computing

The SSP mission, in moving towards predictive simulation, requires (1) improved science-based physical models (which are more computationally intensive), (2) improved geometric fidelity (higher resolution), (3) more common use of 3D simulations in ongoing stockpile work, and (4) very large suites of simulations necessary to support the UQ analyses to provide confidence in the simulation capabilities. The requirement for computational power coming from these four areas is multiplicative and, frankly, daunting. These requirements, however, inform the ASC platform strategy. The largest systems for near-term procurement (within the next three years) are in the 1–20 petaFLOPS (PF) range. Current projections indicate that over the next decade, computational capabilities necessary to meet SSP mission needs will be at the Exascale level, a 1,000-fold to 100,000-fold increase in capability from ASC's largest currently deployed system. The relationship between the development and implementation of improved physics models in the Integrated Design Codes (IDCs) and the need for multi-exaFLOPS (EF) computing resources is summarized in the figure below.

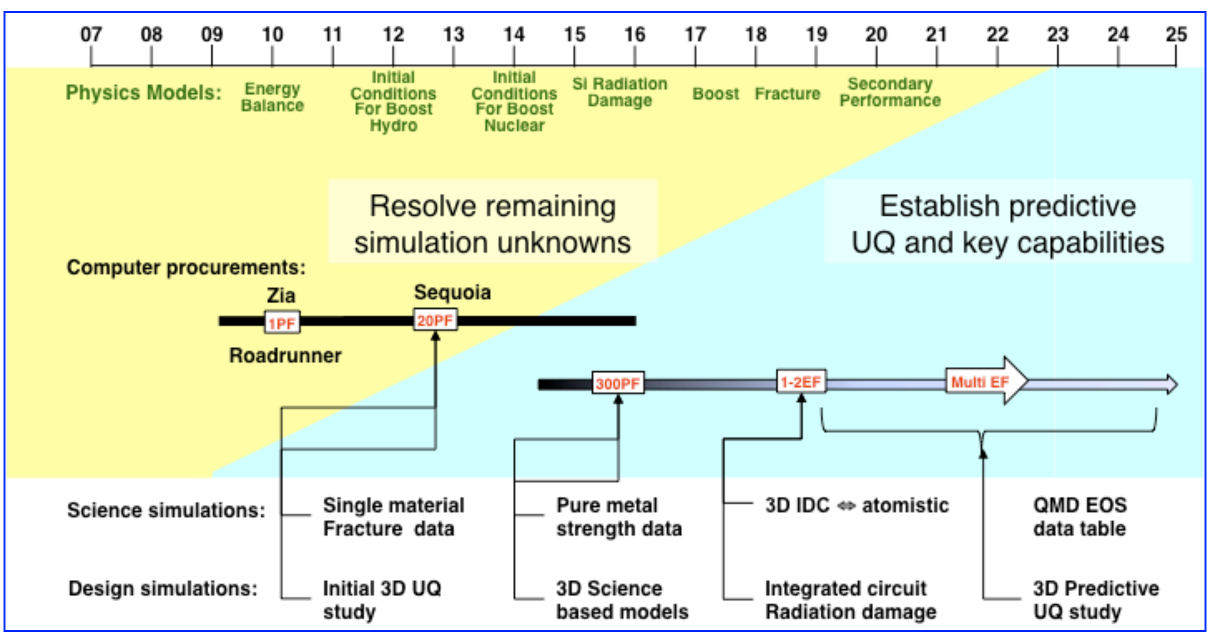

4.2 Delivering Exascale Computing Capability: Platforms, Environment, and Applications Software

As was seen in the preceding figure, in addition to the science base and predictive simulation capability necessary to support future stockpile stewardship, the second challenge to ASC over the next decade will be developing the multi-exaFLOPS computing capability required to satisfy the projected SSP mission needs—the computer hardware platforms, the supporting software infrastructure and, finally, the application codes capable of running on systems architecturally very different than those in use today.

Exascale computing is a capability that has been identified by multiple agencies to address a number of national challenges.[6] NNSA, through ASC, is a leader in HPC and is differentiated from other potential participants through its consistently strong mission focus. Nonetheless, exascale computing is a major technological challenge with a significant requirement for investment in research and development, and effective cooperation between interested parties will be required to achieve this goal in a time frame necessary to meet SSP mission needs.

4.2.1 Building Exascale Computing Platforms

Impressive progress in hardware, software, and tools was made over the first decade of the ASC Program. In simple terms, the basic paradigm to get to where we are today was to connect large numbers of relatively standard computers (processors) via high-speed interconnects into massively parallel clusters. The computing power increased with the number of processors (known as weak scaling) and through the increases in processor capabilities (Moore's law). The programming model for application codes was relatively stable, using parallel programming and communicating between processors with the standardized MPI. Investments in software development were required to move integrated design codes from the older vector-machine-based coding to the newer massively parallel MPI-based designs. Computers got bigger, and each generation had its own challenges, but the basic parallel programming model using MPI remained viable, which provided stability in porting the application codes to each new generation of ASC machines. This paradigm successfully led from the initial ASCI machines to the current ASC petascale capabilities. We are, however, reaching the practical limits of this approach. As the recent DARPA (Defense Advanced Research Projects Agency) announcement noted, "Until recently, advances in Commercial Off-The-Shelf (COTS) systems performances were enabled by increases in clock speed, decreases in supply voltage, and growth in transistor count. These technology trends have reached a performance wall where increasing clock speed results in unacceptably large power increases, and decreasing voltage causes increasing susceptibility to transient and permanent errors. Only increasing transistor count continues to drive performance increases, with value only if we can minimize energy while optimizing our ability to efficiently utilize available concurrency. Further, increasing density has not helped reduce the energy costs of data transport either across a chip, between neighboring chips, or between chips on disparate boards. Current interconnect protocols are beginning to require energy and power budgets that rival or dwarf the cost of doing computation."

Another success of the ASC platform strategy has been the close work with industry and leveraging COTS technology when possible. HPC is not a large enough market by itself to unilaterally affect the major directions of chip manufacturers, but by maintaining leadership in the HPC field, ASC can affect aspects of chip designs that are consistent with the direction industry is moving while reaping the cost benefits of mass-produced products.

The challenge of developing an exascale computing capability calls for a partnership between ASC and Office of Science, which openly engages other government agencies, industry and academia. Such a DOE initiative for open science and national security would assure continued American leadership and competitiveness in this strategic area of simulation, and a partnership would be cost effective and synergistic.

4.2.2 Developing Exascale Computing Software

In view of the significant changes in computer hardware architectures necessary with exascale systems, a new programming model will be a critical component of an initiative to build effective exascale computing systems. With clock speeds projected to be flat or even dropping in order to save energy, performance improvements within a chip will come from increased parallelism. It would be premature to rule out any of the architectural models for increasing on-chip parallelism, yet history suggests that a programming model specialized to a single architecture is doomed to fail. Even if architectures become somewhat specialized to a class of applications, the programming model must be portable across all viable architectures. The exascale software effort therefore needs to allow architectures to pursue multiple hardware solutions, while programming models need to support a range of possible solutions. The evolution of a new programming model should be managed and coordinated by the proposed DOE initiative mentioned earlier.

Such changes in the basic programming model will require renewed investments for software algorithms and tools and for adapting the ASC integrated design codes to new programming models. This is analogous to the re-engineering of the codes required in the 1990s when ASCI moved from vector-based machines to massively parallel clusters programmed with MPI. ASC is additionally constrained by its need to provide continuity of operations. That is, current codes must continue to run to support the ongoing work of the SSP at the same time as efforts are made to accommodate future hardware architectures.

4.2.3 Addressing Exascale Computing Challenges

Given the significant investments already made in ASC integrated design and science codes, ASC must play a central role in the development of exascale hardware and software in order to limit the potential damage to previous investments resulting from suboptimal technologies chosen for future exascale systems. Good choices can be made, and ASC has an unparalleled and exemplary history in making productive choices through deep partnerships with key American HPC vendors. However, achieving exascale computing is a significant technical challenge that is beyond the reach of ASC as it is currently funded; this will most likely require a DOE-wide solution, including effective teaming with industry, academia, Office of Science labs, and other agencies such as DARPA.

DOE laboratories from both NNSA and the Office of Science are actively working together to plan and size a proposed DOE-wide initiative to achieve productive exascale environments and applications over a decade of focused work. In June 2009, this collaborative exascale initiative was publicly announced at the SciDAC conference, and it will be led by the NNSA (ASC) and Office of Science (ASCR). It is a science- and mission-driven DOE initiative to achieve major advances in predictive capability for critical mission areas, including climate, energy, security, and fundamental sciences. An important component of such an effort would be the careful examination of candidate technologies and the pursuit of at least two approaches, thus making the best use of technological diversity in achieving exascale computing while reducing overall risk.

4.3 Meeting Program Challenges within Existing Budget Profiles

The current budget profiles in the Future-Years Nuclear Security Plan (FYNSP) shows funding for ASC relatively flat over the next five years. This reflects NNSA's view of resources necessary to meet current activities. The overall ASC Program budget has experienced gradual but consistent decline in the past several years. The reasons for this decline are multifaceted; however, given the increasing cost of resources to develop code capabilities, to address changes in computing hardware, to develop improved physical models, and to develop UQ as a discipline, at the current levels of funding the program will face difficult decisions regarding which capabilities to sustain, which to neglect, and which to advance.

ASC recently undertook an exercise to estimate, in a budget-independent manner, the staffing needs of the codes and modeling side of the program to support SSP mission requirements. NNSA is extending this exercise to the science and engineering campaigns. The preliminary conclusion is that the current size of the ASC modeling effort has crossed below what is needed to maintain *today's* stockpile and steadily pursue predictive capability. Continuing to maintain the current stockpile while making progress towards a predictive capability and while pursuing exascale computing in a timeframe that supports the SSP will require additional staffing and more effective cross-agency leveraging of resources. Moreover, within a flat overall budget profile and a corresponding platform budget of around $75M/year, ASC alone will be unable to deliver exascale computing capabilities by the end of the next decade as called for in the *ASC Roadmap*.

Delivering a predictive computing capability at the exascale in a time-urgent manner to assure long-term confidence in the nuclear deterrent will require new investments in applications, science, systems hardware and software, and DOE-wide partnerships that are not currently identified in the budgets. Such progress will require a sense of national purpose and significant investments over a decade.

5 ASC Platform Acquisition Strategy

Future SSP challenges are driving the long-term ASC platform plan. ASC's overall platform strategy is described at length in an NNSA document[4] that describes four main principles that guide the acquisition of platforms:

1. Maintain continuity of production.
2. Ensure that the needs of the current and future stockpile are met.
3. Balance investments in system cost-performance types with computational requirements.
4. Partner with industry to introduce new high-end technology constrained by life-cycle costs.

An overview of the ASC platform acquisition strategy is graphically represented in the figure below.

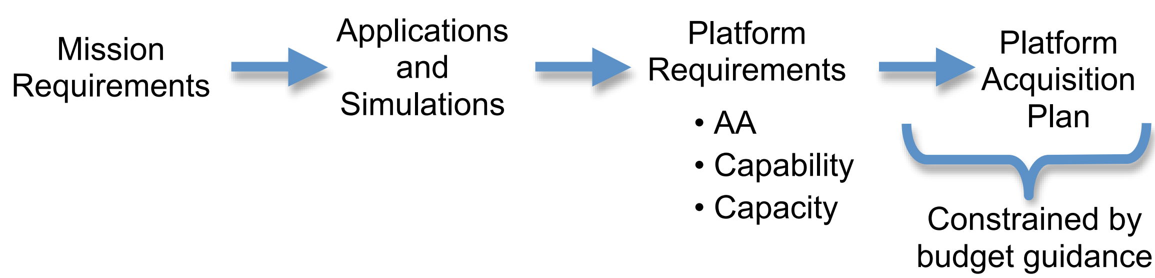

The ASC Program strategy requires the use of a mix of different systems to achieve its mission. In order to mitigate the cost of these systems, ASC distributes the workload onto the most cost-effective computer for the task at hand. Three types of major systems—*capacity*, *capability*, and *advanced architecture (AA)*—are necessary to meet current mission needs and support the development of future capabilities. During any single year, ASC may invest in all three types of computing to various degrees. The amount of investment in each type of system will vary, depending on available resources and mission needs.

Capacity systems generally represent lower-risk, less expensive production computing systems that run parallel problems with more modest computational requirements. Capacity systems are the primary work tool of designers who now routinely work at 1,000+ processors. These systems primarily address the high throughput computing regime referenced in Table 1.1.

Capability systems are general-purpose production systems representing leadership-class machines, dedicated to the most challenging problems of the Nuclear Weapons Complex. These are among the largest systems in the world at a given time, and they have the computational power, memory size, and interconnect speed necessary to address complex weapons system problems. These systems primarily address the high-resolution, multi-physics regime referenced in Table 1.1.

Advanced architecture systems extend the limits of technology by exploring promising architectural approaches. These systems are typically costly because they lie at and test current technological boundaries. A major facet of the AA systems element of the ASC strategy is to enable the program to identify and develop the technology used by future capability and capacity systems and also to engage the major HPC companies in ways that help steer commercial technology to solutions that can be leveraged for national security needs. These systems are also powerful resources for important high- resolution weapons science simulations necessary for improving the physics models as well as for running UQ suites. Consequently, while they test future cost-effective technologies they are simultaneously used for essential programmatic work.

All three types of systems are critical to ASC's ability to support the SSP mission. Prior to the acquisition of any major system, a Critical Decision 0 (CD0) document is prepared and approved that presents the detailed mission needs driving the acquisition. This enables a given acquisition to be prioritized among the various investments that the SSP must make.

Each year the ASC Program evaluates its ability to accomplish its mission and prioritizes its work scope within available resources across a five year horizon. In a separate document that is updated regularly, detailed, *budget-driven* planning for ASC's acquisition of the three types of major systems is laid out over the next five years based on prevailing guidance in the FYNSP. This acquisition plan is currently based on a flat funding profile and will maintain, and to some degree extend, the current computational capability of the program.

Most of the technical issues involved in achieving exascale computing are beyond the scope of the Platform Acquisition Strategy. Such details are being addressed by the recently formed ASC-ASCR exascale collaboration, where the first task is to define the scope and projected costs for achieving exascale computing by the end of the next decade. The work of this collaboration will inform the advanced architecture element of ASC's Platform Acquisition Strategy.

6 Delivering NNSA Computing Capability to the Science-Based Stockpile Stewardship Program

The three NNSA laboratories each have ongoing responsibilities for specific weapon systems, annual assessment of stockpile systems, LEPs, and SFIs, and increasingly important roles to proactively and independently peer review each other's work. Without full-scale nuclear testing to provide an absolute assessment, peer review must be one tool used to fill that void. The specific activities and computing needs at each laboratory change from quarter to quarter. The ASC Program provides computational resources to the SSP necessary to meet these mission needs based on a holistic view of resources available across the tri-laboratory. As a consequence of these various responsibilities, current stockpile computing demands keep existing resources fully utilized, and the demand currently exceeds the available capacity.

Low cost *capacity* systems are provided to all three NNSA laboratories, based on shorter term, less computationally demanding program-driven requirements, to enable flexibility in meeting mission needs. While the ASC Program at DOE HQ determines the allocation of computers to each laboratory, each lab determines the allocation of computing time on its capacity resources to workers locally. ASC capacity systems use a common computing environment at all three laboratories, thus enabling the laboratories to leverage efforts to implement and support the computing environment and permitting the shifting of work on capacity systems between the labs to meet exigent needs.

More expensive and formidable *capability* systems are resources to the SSP that are operated as national user facilities for the three labs. Aggregated capability is a cost-effective approach given the magnitude of the investment in these systems. Capability systems are being sited at two locations: at the Alliance for Computing at Extreme Scales (ACES) in New Mexico, jointly managed by Los Alamos and Sandia National Laboratories, and at the Livermore Computing Center (LC) in California, managed by Lawrence Livermore National Laboratory. The first such facility is the Purple facility at LC; in the coming year when the Zia system goes into production at the ACES consortium, that machine will also take on the responsibility as a NNSA user facility. Having more than one site reduces risks, including those posed by natural catastrophes like fire or earthquake. That these resources are not locally sited at each laboratory is not a significant issue because remote classified computing successfully makes these resources available to all three laboratories. More than a decade of experience in remote classified computing has proven this to be a viable approach. Allocation of time on capability systems is done via fixed time period proposals, called Capability Computing Campaigns (CCC), submitted by all three labs to a tri-lab review process. Allocation of these resources is made based on SSP mission priorities, the need for the capability resource, as well as the likelihood of the proposal being able to achieve its goals. The data in the Appendix of Section 10 demonstrates that all three laboratories are benefiting from the capability resources at remotely sited user facilities.

Advanced architecture systems are provided to develop and explore promising architectural approaches for future general-purpose capability and capacity systems and to support weapons science simulations efforts (at the forefront of computational research) and system assessment. Due to the innovative aspect of these machines, the policies for allocating time on them will vary from machine to machine, however, it is expected that the use of these machines by all three laboratories will be accommodated. For example, the next machine, Sequoia, will be shared across laboratories similar to the capability systems to meet science and UQ demands at the three labs and will be operated with a similar governance model for allocating computing time.

ASC resources are the nearly exclusive source of computing for carrying out major SSP goals, including the annual assessment of stockpile systems, LEPs, resolution of SFIs, as well as supporting experimental campaigns and advanced certification methodologies. The degree of formality in gathering computing requirements for the program varies, but processes and tools are being put into place to increase uniformity. The bulk of the computational cycles for ongoing, time-constrained work on the stockpile are currently provided by the ASC capability and capacity systems.

7 Concluding Remarks

To quote from the *ASC Platform Strategy*[4]

> The platforms component of the ASC Program is faced with making choices among competing priorities and must select from available options with mission goals in mind. The constraints that the Program faces are limited resources to expend on the tri-lab computing infrastructure, minimizing the disruptive element of new architectures that require rewriting codes, and not unnecessarily imposing new programming models that create serious difficulties for code development. The overriding objective is to maximize the productivity of users and developers while at the same time providing the capability to enhance confidence in simulations of device performance outside the data-range provided by the nuclear test base.
>
> The ASC Program, like its ASCI predecessor, recognizes a national responsibility to ensure that the commercial computing sector, for whom this is a small component of their business, continues to pursue technology advances that enable large-scale scientific explorations for both weapons and nonweapons related problems. The national security enterprise understands the need to drive the industry in directions to ensure the specific program-driven resources will be available when needed, and to influence, to the extent possible, new technology directions.

NNSA's ASC Program has developed competencies in secure high-end computing that have no equal in the world, and if it is to be successful in its mission it must continue to push the boundaries of computing, always keeping foremost in mind that it is the national security mission that is, in the end, the reason for its requisite vitality.

8 References

1. *Advanced Simulation & Computing: The Next Ten Years*, Dimitri F. Kusnezov, August 2004, NA-ASC-100R-04-Vol. 1-Rev. 0.

2. *Advanced Simulation & Computing: Business Model*, Dimitri F. Kusnezov, July 2005, NA-ASC-104R-05-Vol. 1-Rev. 0.

3. *Advanced Simulation & Computing Roadmap: National Nuclear Security Through Leadership in Weapons Science*, Dimitri F. Kusnezov and Njema Frazier, November 2006, NA-ASC-105R-06-Vol. 1-Rev. 0.

4. *A Platform Strategy for the Advanced Simulation and Computing Program*, Robert Meisner, August 2007, NA-ASC-113R-07-Vol. 1-Rev. 0.

5. *Advanced Simulation and Computing National Code Strategy: Simulation-Based Complex Transformation*, Kenneth Alvin, Njema Frazier, Robert Meisner, January 2009, NA-ASC-108R-09-Vol. 1-Rev. 0.

6. *ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems*, Peter Kogge *et al*., September 28, 2008, DARPA IPTO/AFRL.

9 Review Panel Members

Dr. Douglas B. Kothe
Director of Science for the National Center for Computational Sciences
Oak Ridge National Laboratory

Dr. Horst D. Simon
Associate Laboratory Director for Computing Sciences
Lawrence Berkeley National Laboratory

Prof. Rick L. Stevens
Associate Laboratory Director for Computing, Environment and Life Sciences
Argonne National Laboratory
Professor of Computer Science and Senior Fellow of the Computation Institute
University of Chicago

10 Appendix:
Data on Capability Computing Allocation and Application to the SSP

This appendix provides data from the ASC capability system, Purple, indicating the allocation of capability computing time to each lab, and also how those allocations are being used to address SSP goals. As mentioned in Section 6, capability resources are accessible to users at all three laboratories, and are obtained through proposals to a tri-lab review process. Each Capability Computing Campaign (CCC) lasts approximately six months. Requests for allocations have consistently exceeded available time by a factor of three since this process was initiated three years ago. Similarly, the CD0 document for the upcoming Zia machine also demonstrates that growing SSP mission needs for capability computing are well in excess of available resources.

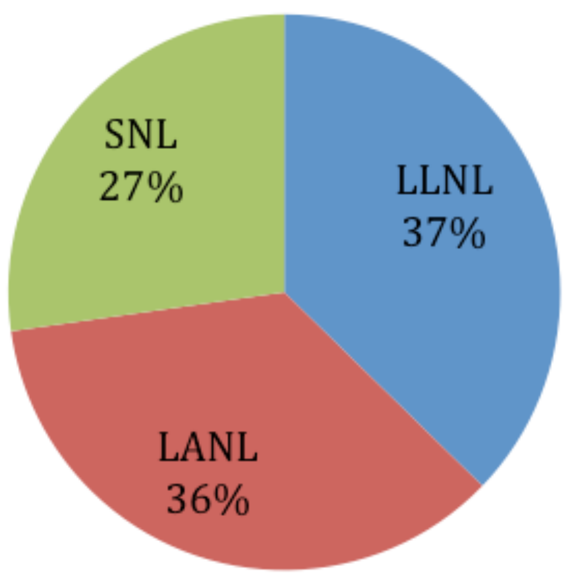

Figure 10.1 – Allocation of cability resources, by laboratory, managed through the competitive CCC process. This represents cumulative data over six CCC campaigns, since Purple began operating as a national tri-lab user facility in October 2006.

Very high system utilization (in excess of 90% over three years*) has been maintained with excellent tri-lab customer feedback on the services provided by Purple. Some representative data is provided below.

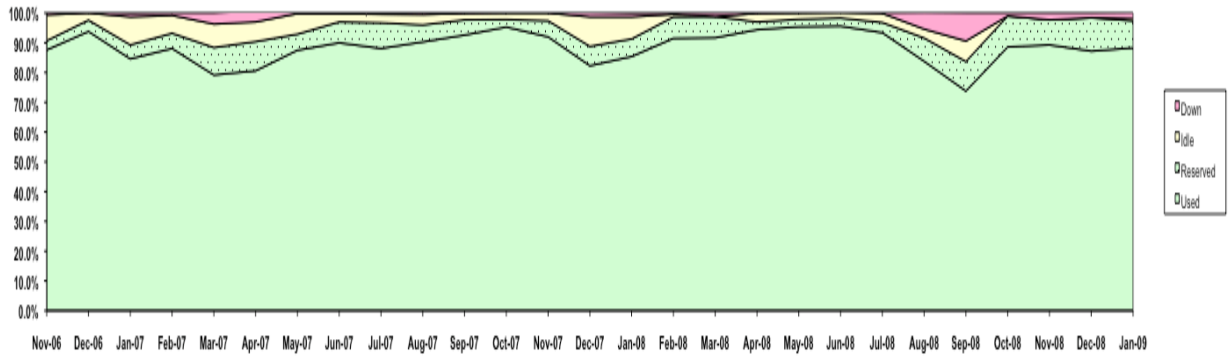

Figure 10.2 – Data showing high utilization of ASC capability resources (Purple).

The SSP goals addressed by work packages vary from proposal to proposal and from campaign to campaign. Each CCC has delivered important programmatic results for the SSP program. For DSW, capability computing has been used for the W76 LEP, the B61, the B83, the W80 and other systems. It has supported Z, NIF (target performance and laser plasma interactions) and DAHRT. In terms of advancing predictive capability, it has heavily impacted both Boost and Energy Balance. In terms of the former, it was phenomena discovered at exquisite resolution on Purple that was the catalyst for the National Boost Initiative and indeed the quest for predictive simulation. In addition, Purple completed some full system (8,000 core) calculations that will be essential to the resolution of the Energy Balance physics model. As a representative example, the table below shows how work on the current campaign is allocated across specific SSP mission areas.

SSP Goals		
DSW (LEPs, SFIs, etc.)	Advance Predictive Capability	Support and Design of Experiments
44%	28%	27%

Table 10.3 – Data from the most recent Capability Computing Campaign showing the SSP mission areas addressed by the work packages selected for execution.

* Utilization as tracked by NNSA includes the sum of the time jobs are running on nodes and the time that nodes are held in reserve in order to create enough "space" for the next job waiting. This latter category is called "reserved" and is typically about 7.5%. Consequently, utilization of 97% means about 90% is used for jobs, about 7% is reserved and 3% is either down or idle. Purple utilization statistics are considered extremely high.